RCA 89648

CITED BY APPLICANT

VSB MODEM SUBSYSTEM DESIGN FOR GRAND ALLIANCE DIGITAL TELEVISION RECEIVERS

Wayne Bretl, Gary Sgrignoli, and Paul Snopko
Zenith Electronics
1000 Milwaukee Ave.
Glenview, IL 60025

Abstract: The functions of the vestigial sideband modem for the Grand Alliance digital TV are described. The partitioning of these functions is discussed in terms of the prototype hardware and possible future development.

Summary
The vestigial sideband (VSB) modem subsystem (tuner/demodulator) for the Grand Alliance digital TV receiver includes the functions of RF selection, local oscillator, and conversion; IF amplification and band shaping; baseband demodulation; AGC; co-channel filtering; synchronization and phase tracking of the carrier, the bit clock, and the data framing; equalization (ghost canceling); forward error correction (trellis and RS codes); and data de-interleaving. (Figure 1) The interaction and partitioning of these functions is discussed.

RF, Local Oscillator And Conversion

The prototype hardware uses a double-conversion system, with high-side injection and the first IF at 920 MHz, to tune the entire VHF, UHF and cable bands. This places image frequencies above 1 GHz, where they are easily rejected by an input filter. (See Figure 2).

The input filter is a broadband bandpass 50-810 MHz tracking filter. This tracking filter is not narrow or critically tuned, as it must be in present NTSC tuners which must reject signals only 90 MHz away from the selected channel.

The 978-1723 MHz first local oscillator (LO) frequency also is prevented from leaking out of the input by the bandpass filter. At the same time, second harmonics of UHF channels (2x470 to 2x806 MHz) fall above the first IF bandpass. Harmonics of cable channels could possibly occur in the first IF passband but are not a problem because of the relatively flat spectrum (within 10 dB) and small signal levels (-28 dBm or less) used in cable systems.

Figure 1

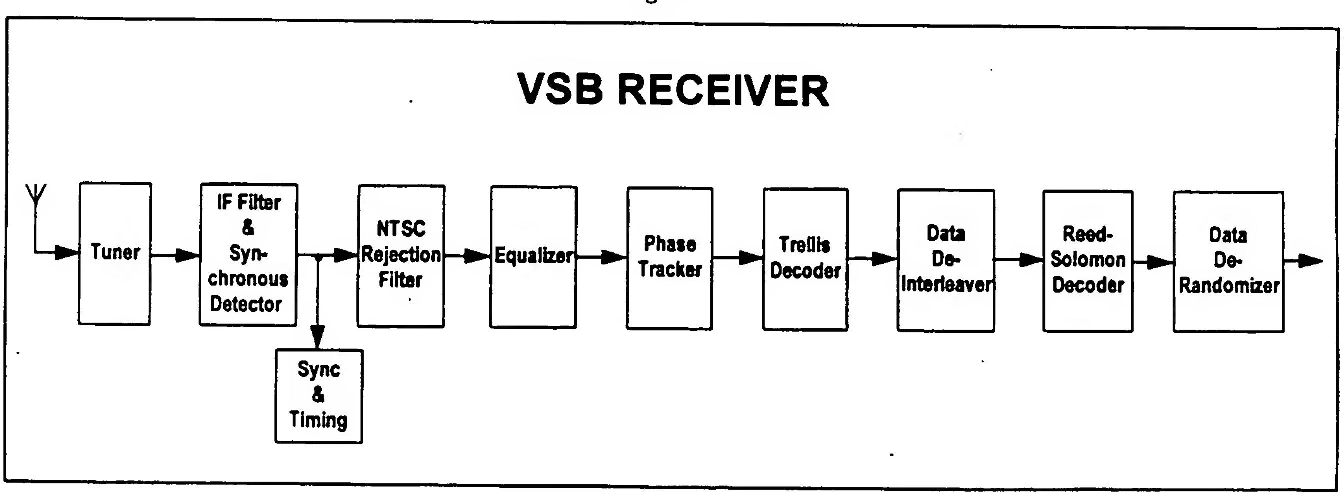

Manuscript received June 12, 1995

Figure 2

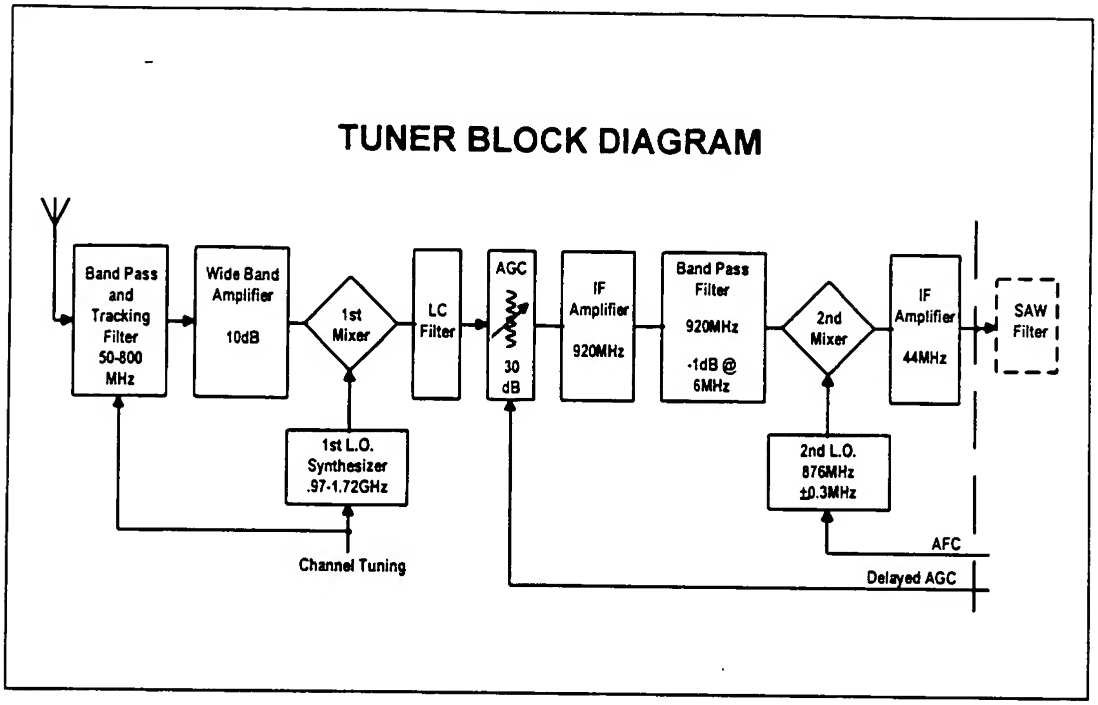

A first RF stage with 10 dB gain essentially determines the receiver noise figure, 7 to 9 dB over the entire tuning range. The prototype hardware uses a highly linear double balanced first mixer design to minimize even harmonic distortion. It is driven by a synthesized low phase noise first LO above the first IF frequency (high-side injection). Both the channel tuning (first LO) and broadband tracking filtering (input bandpass filter) are controlled by a microprocessor. This arrangement is capable of tuning the entire VHF and UHF broadcast bands as well as all standard, IRC, and HRC cable bands.

The mixer is followed by an LC filter in tandem with a narrow 920 MHz ceramic resonator bandpass filter. The LC filter provides selectivity against the harmonic and subharmonic spurious responses of the ceramic resonators. The 920 MHz ceramic resonator bandpass filter has a -1 dB bandwidth of about 6 MHz. A 920 MHz IF amplifier is placed between the two filters. Delayed AGC of the first IF signal is applied immediately following the first LC filter. The 30 dB range AGC circuit protects the remaining active stages from large signal overload, while the second IF provides additional AGC action in a stage following the SAW filter.

The second mixer is driven by the second LO, which is an 876 MHz voltage-controlled SAW oscillator. It is controlled by the frequency and phase-locked loop (FPLL) synchronous detector (Figure 3). The second mixer, whose output is the desired 44 MHz second IF frequency, drives a constant gain 44 MHz amplifier. The output of the tuner feeds the IF SAW filter and synchronous detection circuitry.

The VSB tuner system can meet phase noise and other requirements with standard consumer components and stamped-metal construction.

VSB Carrier Recovery

Carrier recovery is performed by a frequency- and phase-lock loop (FPLL, Figure 3) using the small pilot carrier included in the VSB signal. The first LO is synthesized by a PLL and controlled by a microprocessor. The third LO is a fixed reference oscillator. Any frequency drift or deviation from nominal has to be compensated in the second LO.

Control for the second LO comes from the FPLL synchronous detector, which integrally contains both a frequency loop and a phase-locked loop in one circuit. The frequency loop provides a wide frequency pull-in range of ±100 kHz while the phase-locked loop has a narrow bandwidth (less than 2 kHz).

Figure 3

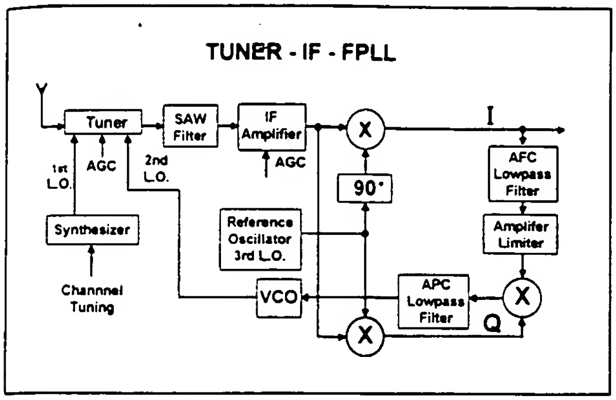

During frequency acquisition, the FPLL uses both the in-phase (I) and quadrature-phase (Q) pilot signals. All other data processing circuits in the receiver use only the I-channel signal. Prior to phase-lock, as is the condition after a channel change, the automatic frequency control (AFC) lowpass filter acts on the beat signal created by the frequency difference between the VCO and the incoming pilot. The high frequency data (as well as noise and interference) is mostly rejected by the AFC filter, leaving only the pilot beat frequency. The pilot beat signal is limited to a constant amplitude (but maintaining its positive or negative polarity), and is used to multiply the quadrature signal, resulting in a traditional bipolar S-curve AFC characteristic. The polarity of the S-curve error signal depends upon whether the VCO frequency is above or below the incoming IF signal. Filtered and integrated by the automatic phase control (APC) lowpass filter, this DC signal adjusts the tuner's second LO to reduce the frequency difference.

When the frequency difference comes close to zero, the FPLL reverts to phase control and phase-locks the incoming IF signal to the third LO. This is a normal phase-locked loop circuit, with the exception that it is bi-phase stable. However, the correct phase-lock polarity is determined by forcing the polarity of the pilot to be equal to the known transmitted positive polarity. Once locked, the detected pilot signal is constant, the limiter output feeding the third multiplier is at a constant positive polarity, and only the phase-locked loop is active (frequency loop automatically disabled). The APC lowpass filter is wide enough to allow reliable ±100 kHz frequency pull-in, yet narrow enough to reject consistently all strong white noise and NTSC cochannel interference. The PLL has a bandwidth that is narrow enough to reject also most of the AM and PM generated by the random data, yet is wide enough to track out any phase noise on the signal (and, hence, on the pilot) out to about 1 kHz. Tracking out low frequency phase noise (as well as low frequency FM components) allows the following phase tracking loop, discussed later, to be more effective.

The prototype receiver can acquire a signal and maintain lock at a signal-to-noise ratio of 0 dB or less, and in the presence of heavy interference.

Segment Sync and Symbol Clock Recovery

Data segment sync in the terrestrial and cable signals are shown In Figure 4 and Figure 5 respectively. The repetitive data segment syncs are separated from the random data by a narrow bandwidth segment integrator circuit. The data segment syncs are used to regenerate a properly phased 10.76 MHz symbol clock and can also be used to generate a coherent AGC control signal. The segment sync amplitude and pilot amplitude are expressed in terms of the integer-numbered data levels in each case.

A block diagram of the circuit and details of the sync recovery are shown in Figure 6. The 10.76 Msymbols/sec I-channel baseband signal (syncs and data) from the synchronous detector is converted by an A/D converter for digital processing. Traditional analog data eyes can be viewed after synchronous detection. However, after conversion to a digital signal by T sampling, only the value of the signal at the sampling instant is retained

The 4-symbol correlation filter detects the two-level syncs (and any data patterns that resemble 2-level sync). This always occurs, irrespective of whether the voltage-controlled crystal oscillator (VCXO) is locked or free-running. Repeated correlations at a data segment rate are integrated and amplified by the segment integrator. Thus, the repetitive segment sync is detected, while the random data integrates to zero, enabling the loop to lock on the sampled sync only. Precise phasing is obtained by sampling the zero crossing of the quadrature filter output, which corresponds to the center of the sync waveform. Upon reaching a predefined level of confidence that the segment sync has been found (using a confidence counter), subsequent receiver loops are enabled.

Data segment sync detection and clock recovery both work reliably at signal-to-noise ratios of 0 dB or less, and in the presence of heavy interference.

Figure 4

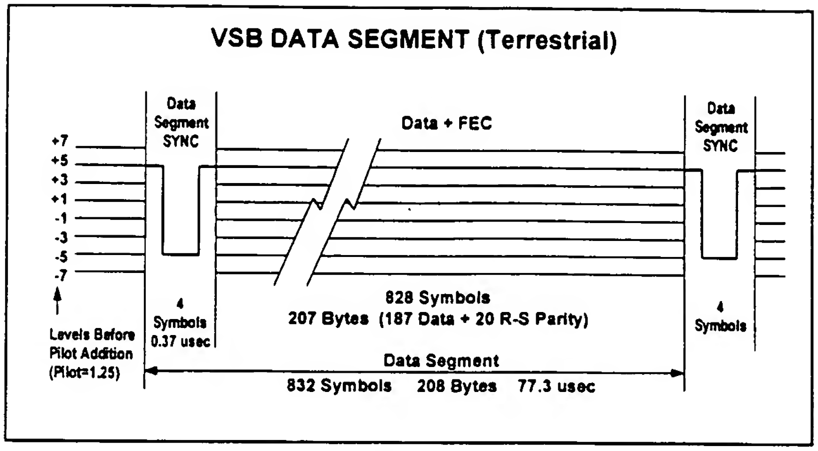

Figure 5

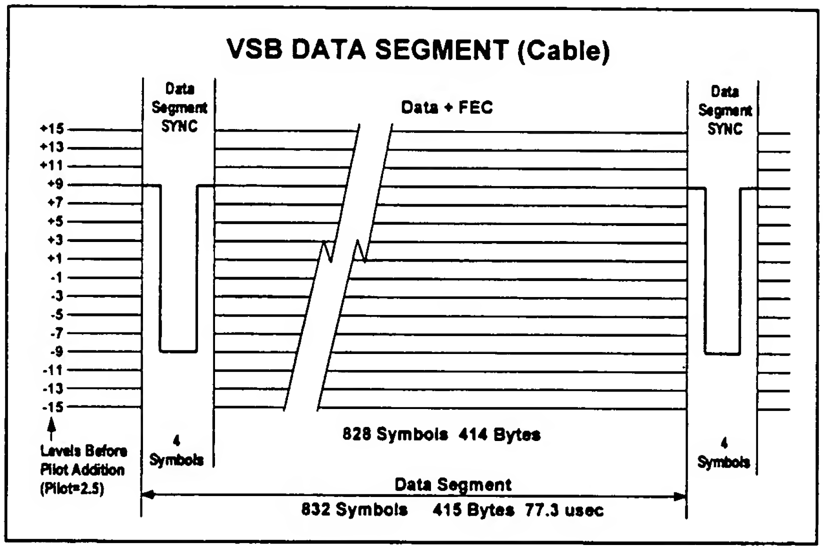

Figure 6

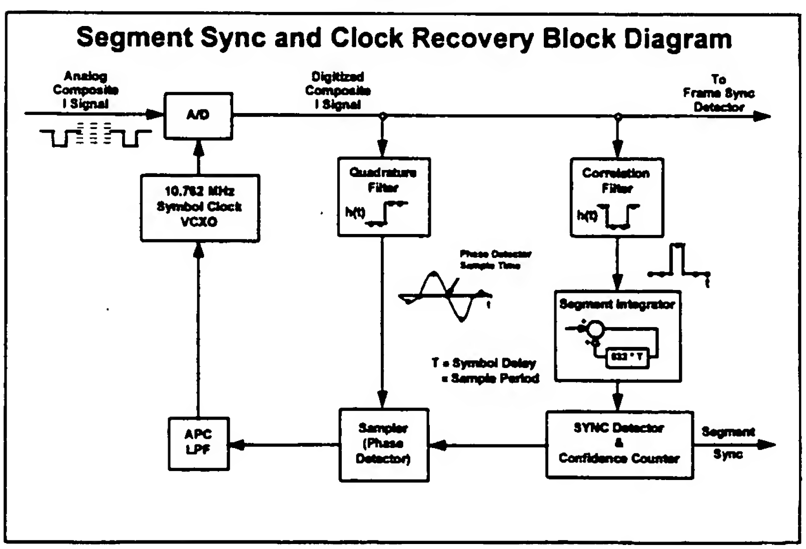

Non-Coherent and Coherent AGC

Prior to carrier and clock synchronization, non-coherent automatic gain control (AGC) is performed on any signal (locked or unlocked signal, or noise/interference) that is received. The IF and RF gains are reduced accordingly, with the appropriate AGC "delay" applied.

When data segment syncs are detected, coherent AGC may be applied using the measured segment sync amplitudes. The amplitude of the bipolar syncs, relative to the discrete levels of the random data, is determined in the transmitter. Once the syncs are detected in the receiver, they are compared to a reference value, with the difference (error) integrated. The integrator output then controls the IF and "delayed" RF gains, forcing them to whatever values provide the correct sync amplitudes. Because of the data randomization used in the VSB signal, it is also generally possible to operate AGC on a non-coherent basis.

Data Frame Synchronization

The data frame sync signal is shown in Figure 7. The frame sync is a binary signal of the same amplitude as the data segment sync. Data frame sync detection, shown in Figure 8, is achieved by comparing every received data segment from the A/D converter (after interference rejection filtering to minimize cochannel interference), to an ideal frame reference signal in the receiver.

Oversampling of the frame sync is not necessary since a precision data segment and symbol clock has already been created reliably by the clock recovery circuit. Therefore, the frame sync recovery circuit samples exactly where a valid frame sync correlation should occur within each data segment, and only needs to perform a symbol by symbol difference. Upon reaching a predetermined level of confidence (using a confidence counter) that frame syncs have been detected on given data segments, the Data Frame Sync signal becomes available for use by subsequent circuits. The detection procedure makes frame sync recovery robust even in heavy noise, interference, or ghost conditions. Frame sync recovery can occur reliably at signal-to-noise ratios of 0 dB or less, and in the presence of heavy interference.

In addition to frame synchronization, the frame sync signal provides the training signal for data equalization, and an identification of the VSB mode for establishing the proper slicing levels. (Note that alternating frames have alternating polarities of the three 63-symbol pseudo random "PN63" sequences.)

Figure 7

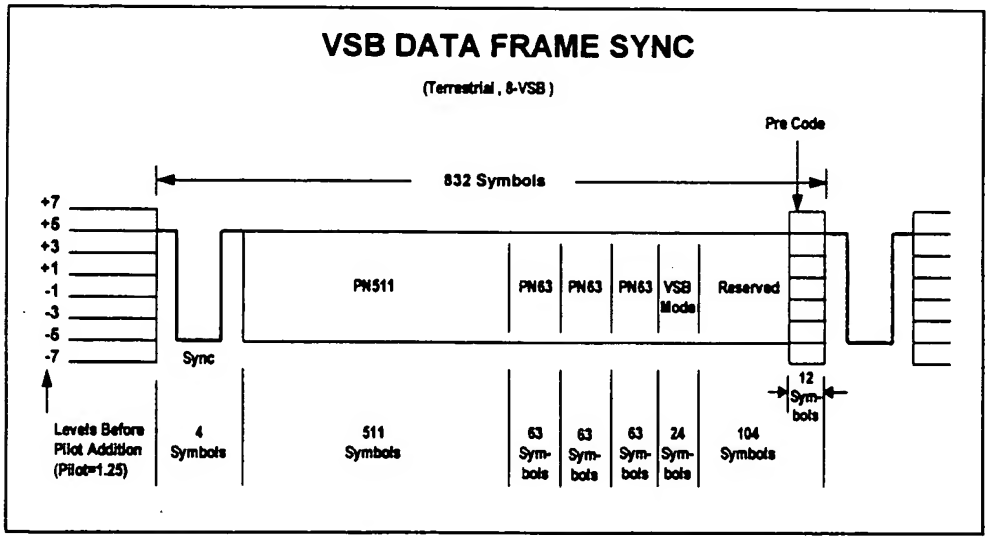

Figure 8

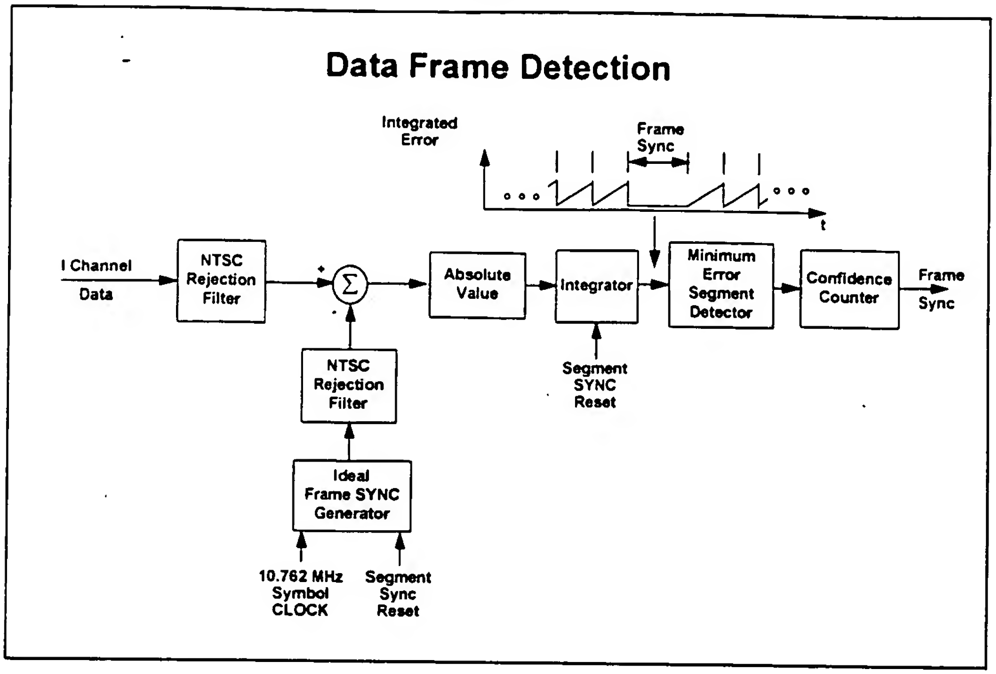

NTSC Co-Channel Rejection

The terrestrial VSB signal is digitally precoded so that it can be received in two modes: directly as an 8-level signal, with digital decoding, or through a NTSC rejection filter (12-symbol comb) which translates it into a 15-level signal. The filter can be switched out, for a 3-dB white-noise-only improvement. The receiver continuously measures the error performance before and after filtering, and automatically uses the best mode according to prevailing noise and/or co-channel conditions. Switching does not interrupt reception, and therefore can occur automatically as conditions change.

Interference Rejection Filter

The interference rejection properties of the VSB transmission system are based on the frequency location of the principal components of the NTSC cochannel interfering signal within the 6 MHz TV channel and the periodic nulls of a VSB receiver baseband comb filter.

Figure 9 shows the location and approximate magnitude of the three principal NTSC components: (1) the visual carrier located 1.25 MHz from the lower band edge, (2) the chroma subcarrier located 3.58 MHz higher than the visual carrier frequency, and (3) the aural carrier located 4.5 MHz above the visual carrier frequency.

The NTSC interference rejection filter (comb) is a one tap linear feed-forward filter, as shown in Figure 10.

Figure 9 shows the frequency response of the comb filter, which provides periodic spectral nulls spaced 57 x f_h (10.762 MHz/12, or 896.85 kHz) apart. (f_h is the NTSC line rate.) There are 7 nulls within the 6 MHz channel. The NTSC visual carrier frequency falls close to the second null from the lower band edge. The 6th null from the lower band edge is nominally placed for the NTSC chroma subcarrier, and the 7th null from the lower band edge is near the NTSC aural carrier.

The visual carrier falls 15.1 kHz above the second comb filter null, the chroma subcarrier falls 7.2 kHz above 6th null, and the aural carrier falls 30.8 kHz above the 7th null. (Note, the aural carrier is at least 7 dB below its visual carrier).

The comb filter, while providing rejection of steady-state signals located at the null frequencies, has a finite response time of 12 symbols (1.115 μsec). Thus, if the NTSC interfering signal has a sudden step in carrier level (low to high or high to low), one cycle of the beat frequency (offset) between the ATV and NTSC carrier frequencies will pass through the comb filter at an amplitude proportional to the NTSC step size as instantaneous interference.

Figure 9

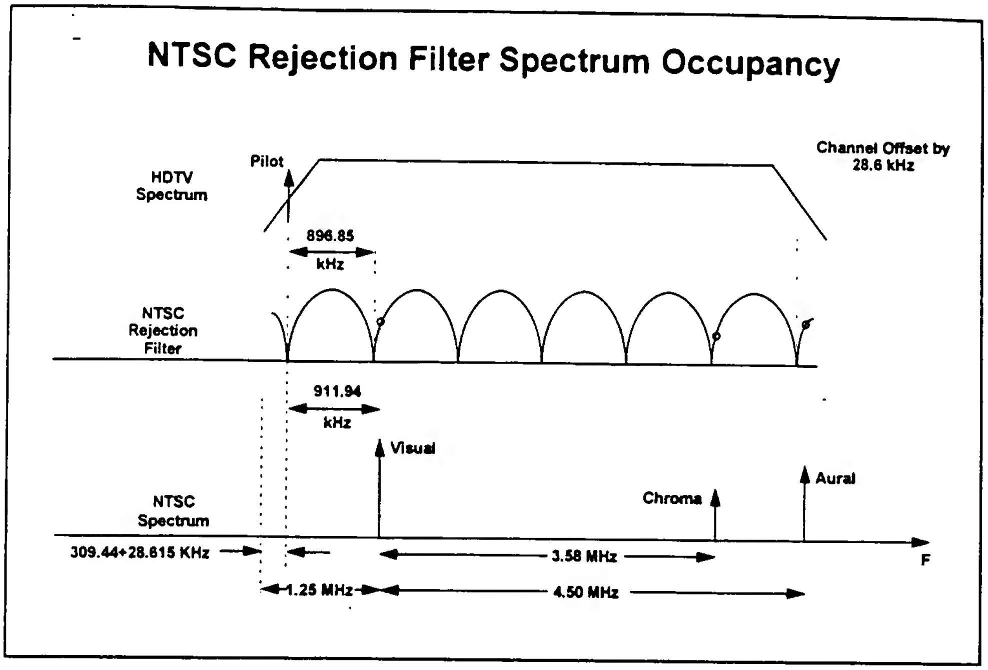

Examples of such steps of NTSC carrier are: leading and trailing edge of sync (40 IRE units). If the undesired signal power is large enough, data errors will occur. However, interleaving will spread the interference and will make it easier for the Reed-Solomon (R-S)code to correct the errors. (R-S can correct up to 10 byte errors/segment).

Although the comb filter reduces the NTSC interference, the data is also modified. The 7 data eyes (8 levels) are converted to 14 data eyes (15 levels). This doubling of the eyes is caused by the partial response process which is a special case of inter-symbol interference that does not close the data eye but creates double the number of eyes of the same magnitude. The modified data signal can be properly decoded by the trellis decoder described in later sections. Note that because of time sampling, only the maximum data eye value is seen after A/D conversion.

Figure 9 and Figure 19 show that there is a shift of the actual VSB carrier frequency compared to the nominal 6 MHz channel. (While not particularly critical, this shift affects interference performance slightly and may be adjusted in the final FCC rules as the result of official tests. The shift is negligible compared to the receiver carrier pull-in capability.) The shift equals +28.615 kHz, or about +0.48%. This is slightly higher than channel offsets currently applied to NTSC (+10, 0, or -10 kHz). It results in the nominal 309.44 kHz roll-off region energy reaching into the upper adjacent channel at a level of about -40 dB or less. If that is another ATV channel, its spectrum is also shifted upward; therefore no spectral overlapping occurs. If it is an NTSC channel, the shift is below the (RF equivalent of the) Nyquist slope of an NTSC receiver where there is high attenuation, and it is slightly above its customary lower adjacent channel sound trap. No adverse effects of the shift have been found nor are they foreseen. An additional shift of the ATV spectrum is used in order to track any dominant NTSC interferor which may be assigned an offset of -10 kHz, 0 kHz, or +10 kHz.

NTSC interference can be detected by the circuit shown in Figure 10, where the signal-to-interference plus noise ratio of the binary Data Frame Sync is measured at the input and output of the comb filter, and these measurements are compared to each other. This is accomplished by creating two error signals. The first is created by comparing the received signal with a stored reference of the frame sync. The second is created by comparing the rejection filter output with a combed version of the internally stored reference frame sync. The errors are squared and integrated. After a predetermined level of confidence is achieved, the path with the largest signal-to-noise ratio (lowest interference energy) is switched into the system automatically.

There is a reason to not leave the rejection comb filter switched in all the time. The comb filter, while

providing needed cochannel interference benefits, reduces white noise performance by 3 dB. This is because the filter output is the difference of two full gain paths, and since white noise is uncorrelated from symbol to symbol, the noise power doubles. There is an additional 0.5 dB degradation due to the 12 symbol differential coding. (See below under Trellis Decoder.) If little or no NTSC interference is present, the comb filter is automatically switched out of the data path. When NTSC service has been phased out, the comb filter may be omitted from ATV receivers.

Figure 10

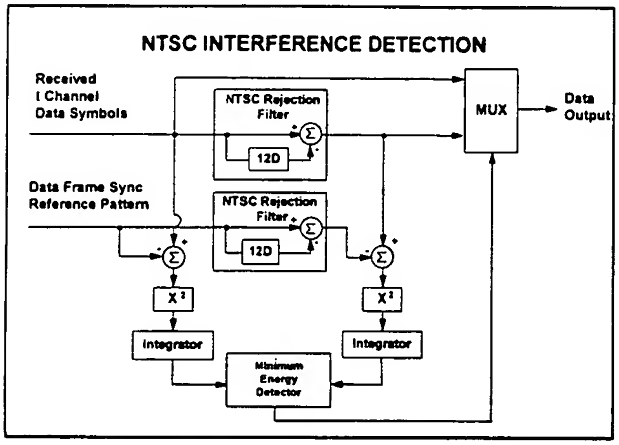

Channel Equalizer / Ghost Cancellation

The equalizer operates on the pseudo-noise (PN) sequence transmitted regularly as data "frame sync" (see Figure 7). In addition to the highly accurate correction possible by this method, the receiver uses data directed equalization for fast-changing conditions, and can use a blind equalization procedure to enhance initial acquisition.

The equalizer/ghost canceller compensates for linear channel distortions, such as tilt and ghosts. These distortions can come from the transmission channel or from imperfect components within the receiver.

The equalizer algorithm can work using three methods: it can adapt on the binary training sequence; it can adapt on data symbols throughout the frame, when the eyes are open; or it can adapt on data when the eyes are closed (blind equalization). The principle difference among these three methods is how the error estimate is generated.

The training signal in the data frame sync presents a fixed repetitive data pattern in the data stream (Figure 11). Since the data pattern is known, the exact error is generated by subtracting the known training sequence from the detector output (see Figure 12

Figure 11

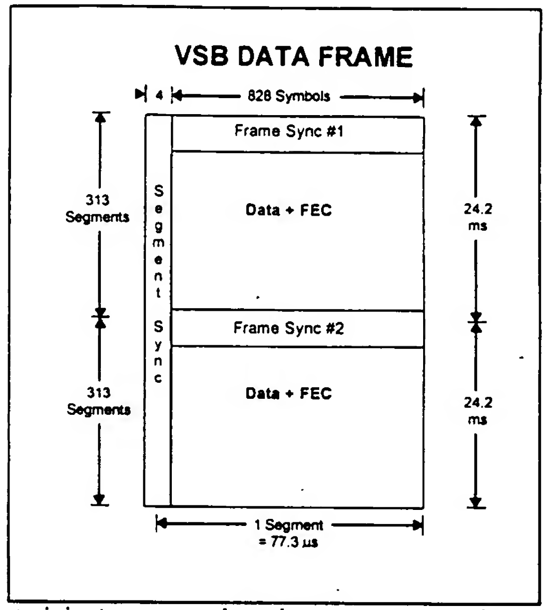

The training sequence alone, however, may not be enough to track dynamic ghosts since these require tap adjustments more often than the training sequence is transmitted. Therefore, once equalization is achieved, the equalizer can switch to adapting on data symbols throughout the frame, and produce an accurate error estimate by slicing the data with an 8-level slicer and subtracting the raw data from the sliced version.

A blind equalization mode can be used for fast initial acquisition (even before the training signal has been used). The blind equalizer models the multilevel signal as a binary data signal plus noise, and the equalizer produces the error estimate by detecting the sign of the output signal and subtracting a (scaled) binary signal from the output to generate the error estimate.

The equalizer uses a least-mean-square (LMS) algorithm. The LMS algorithm computes how to adjust the filter taps in order to reduce the error present at the output of the equalizer. An error estimate (produced using the training sequence, 8-level slicer, or the binary slicer) is multiplied by delayed copies of the signal. The delay depends upon which tap of the filter is being updated. This multiplication produces a cross-correlation between the error signal and the data signal. The size of the correlation corresponds to the amplitude of the residual ghost present at the output of the equalizer and indicates how to adjust the tap to reduce the error at the output.

Figure 12

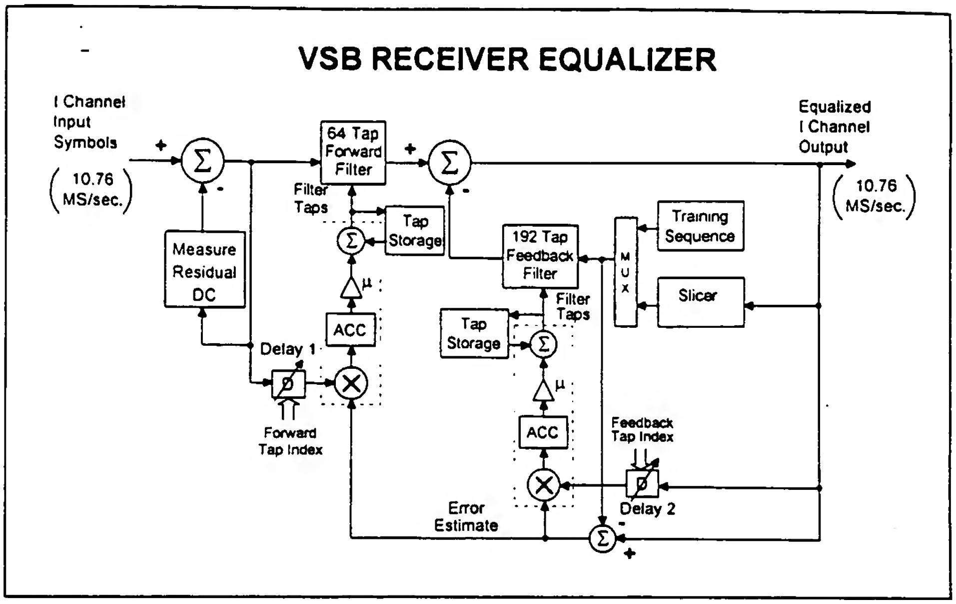

A block diagram of the equalizer is shown in Figure 12. The DC bias of the input signal is first removed by subtraction. The DC may be caused by circuit offsets, non-linearities, or shifts in the pilot caused by ghosts. The DC offset is detected by measuring the DC value of the training signal.

The equalizer filter consists of two parts: a 64 tap feed forward transversal filter followed by a 192 tap decision feedback filter. The equalizer operates at the 10.762 MHz symbol rate (T-sampled equalizer). Note that since VSB modulation is all on one axis, and since accurate symbol timing is available from the segment-sync controlled symbol clock, T/2 or higher sampling is not required.

The output of the feed forward filter and feedback filter are summed to produce the final output. This output is sliced by either an 8-level slicer (15 level slicer when the comb filter is used) or a binary slicer depending upon whether the data eyes are open or not. (As noted in the previous section on interference filtering, the comb filter does not close the data eyes but creates twice as many of the same magnitude). A locally generated training signal and segment syncs are reinserted into the sliced signal (at the "MUX" in Figure 12) since these are known fixed patterns which can be used for equalization error signal generation even when the sliced payload data is corrupted. The resultant signal is fed into the feedback filter, and subtracted from the output signal to produce the error estimate. The error estimate is correlated with the input signal (for the forward filter), or with the output signal (for the feedback filter). This correlation is scaled by a step size parameter, μ, and used to adjust the value of the tap. The delay setting of the adjustable delay is controlled according to the index of the filter tap that is being adjusted.

In addition, the frame sync provides a means to determine conditions of fast-changing ghosts, which requires the data-directed algorithm, or noise, which favors the frame-sync-directed algorithm.

Phase Tracking Loop

The phase tracking loop is an additional decision feedback loop which further tracks out phase noise which has not been removed by the IF PLL operating on the pilot. Thus, phase noise is tracked out by not just one loop, but two concatenated loops. Because the system is already frequency-locked to the pilot by the IF PLL (independent of the data), the phase tracking loop performance is maximized by using a first order loop. Higher order loops, which are needed for frequency tracking, do not perform phase tracking as well as first order loops. Therefore, they are not used in the phase tracker.

A block diagram of the phase tracking loop is shown in Figure 13. The output of the "real" equalizer operating on the I signal is first gain controlled by a multiplier and then fed into a filter which recreates an approximation of the Q signal (Q'). This is possible because of the VSB

transmission method, where the I and Q components are related by a filter function which is almost a Hilbert transform. The complexity of this filter is minor since it is a finite impulse response (FIR) filter with fixed anti-symmetric coefficients and with every other coefficient equal to zero. In addition, many filter coefficients are related by powers of two, thus simplifying the hardware design.

Figure 13

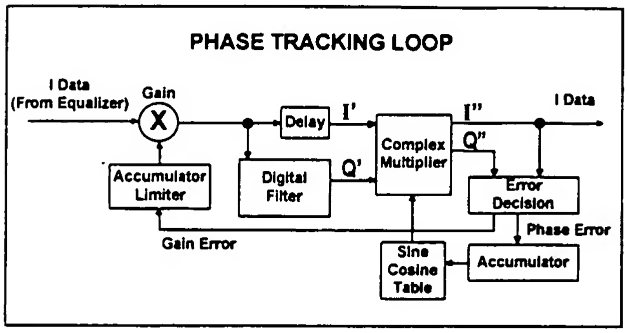

These I' and Q' signals are then fed into a de-rotator (complex multiplier), which is used to remove the phase noise. The amount of de-rotation is controlled by decision feedback of the data taken from the output of the de-rotator. Since the phase tracker is operating on the 10.76 Msymbol/sec data, the bandwidth of the phase tracking loop is fairly large, approximately 60 kHz. The gain multiplier is also controlled with decision feedback.

Forward Error Correction (FEC)

Terrestrial signals (8-VSB) use an inner trellis code to improve white noise performance. Both 8-VSB and 16-VSB use an outer T=10 Reed-Solomon code. In addition, the data is multiplied by a pseudo-random sequence to preserve a random flat spectrum in case the input data is all zero. The receiver corrects and derandomizes the payload data only, because FEC and randomization are not applied to the data segment sync or frame sync signals.

Trellis Decoder

To help protect the trellis decoder against short burst interference, such as impulse noise or NTSC cochannel interference, 12 symbol intrasegment code interleaving is applied in the transmitter. As shown in Figure 14, the receiver uses 12 trellis decoders in parallel, where each trellis decoder sees every 12th symbol. This code interleaving has all the burst noise benefits of a 12 symbol interleaver, but also minimizes the resulting code expansion (and hardware) when the NTSC rejection comb filter is active.

Figure 14

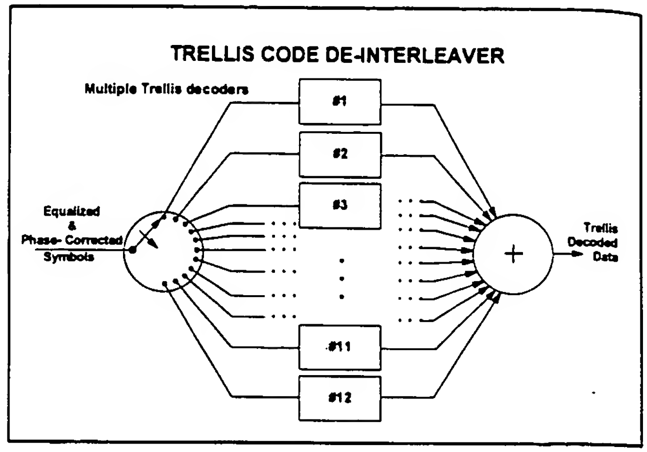

The trellis decoder does slicing and convolutional decoding. It has two modes; one when the NTSC rejection filter is used to minimize NTSC cochannel, and the other when it is not used. This is illustrated in Figure 15. The insertion of the NTSC rejection filter is determined automatically (before the equalizer), with this information passed to the trellis decoder. When there is little or no NTSC cochannel interference, the NTSC rejection filter is not used, and an optimal trellis decoder is used to decode the 4-state trellis-encoded data. Serial bits are re-created in the same order in which they were created in the encoder.

Figure 15

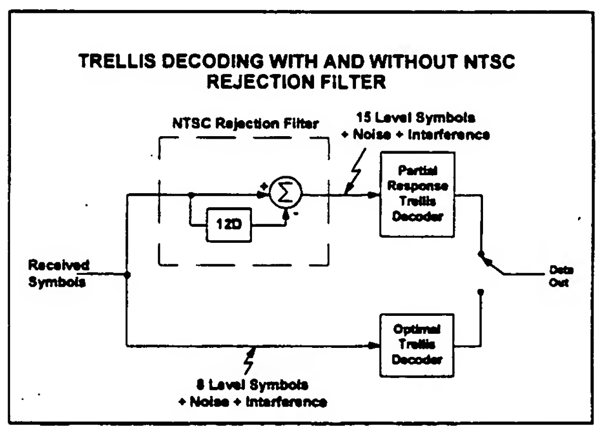

In the presence of significant NTSC cochannel interference, when the NTSC rejection filter (12 symbol, feedforward subtractive comb) is employed, a trellis decoder optimized for this partial response channel is used. This optimal decoder requires 8 states. This is necessary since the NTSC rejection filter, which has

memory, represents another state machine seen at the input of the trellis decoder. In order to minimize the expansion of trellis states, two measures are taken: (1) special design of the trellis code, and (2) twelve-to-one interleaving of the trellis encoding. The interleaving, which corresponds exactly to the twelve-symbol delay in the NTSC rejection filter, causes each trellis decoder to see an apparent NTSC rejection filter that has only a one-symbol delay. By minimizing the delay stages seen by each trellis decoder, the expansion of states is also minimized. A small (3.5 dB) penalty in white noise performance is encountered when the comb filter is used for optimum NTSC cochannel performance. The comb filter itself contributes 3 dB, and an additional 0.5 dB is due to the 12 symbol differential coding.

The presence of the segment sync character in the data stream passing through the comb filter presents a complication which must be dealt with because segment sync is not trellis encoded or pre-coded. Figure 16 shows the technique that has been used. It shows the receiver processing that is performed when the comb filter is present in the receiver. The multiplexer in the Segment Sync Removal block is normally in the upper position. This presents data that has been filtered by the comb to the trellis decoder. However, because of the presence of the sync character in the data stream, the multiplexer selects its lower input during the four symbols that occur twelve symbols after the segment sync. The effect of this sync removal is to present to the trellis decoder a signal that consists of only the subtraction of two adjacent data symbols that come from the same trellis encoder, one transmitted before, and one after the segment sync. The interference introduced by the segment sync symbol is removed in this process, and the overall channel response seen by the trellis decoder is that of the single-delay partial response filter.

Figure 16

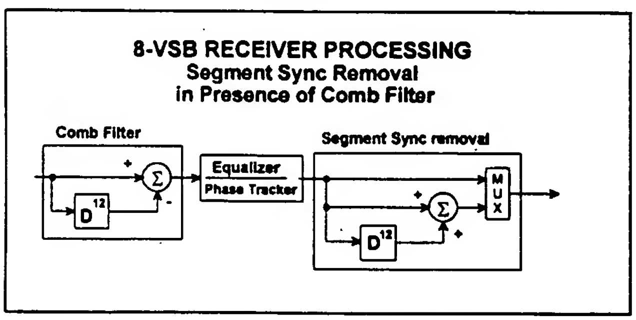

The complexity of the trellis decoder is dependent upon the number of states in the decoder trellis. The maximum requirement is to operate on an 8-state trellis (when the comb filter is active). This defines the maximum processing power that is required. The decoder must perform an add-compare-select (ACS) operation for each state of the decoder. This means that the decoder performs 8 ACS operations per symbol time. When the comb filter is not activated, the decoder operates on a 4-state trellis. The decoder hardware can be constructed such that the same hardware that decodes the 8-state comb filter trellis can also decode the 4-state trellis when the comb filter is disengaged, so there is no need for separate decoders for the two modes. The 8-state trellis decoder requires less than 5000 gates.

It should be noted that after the ATV transition period, when NTSC is no longer being transmitted, the NTSC rejection filter and the 8-state trellis decoder can be eliminated from receivers.

Data De-Interleaver

The convolutional de-interleaver performs the exact inverse function of the transmitter convolutional interleaver. Its 1/6 data frame depth, and intersegment "dispersion" properties allow noise bursts lasting about 193 μsec to be handled. Even strong NTSC cochannel signals passing through the NTSC rejection filter and creating short bursts due to NTSC vertical edges, are handled reliably due to the interleaving and R-S coding process. The de-interleaver uses data frame sync for synchronizing to the first data byte of the data frame. The convolutional de-interleaver is shown in Figure 17.

Figure 17

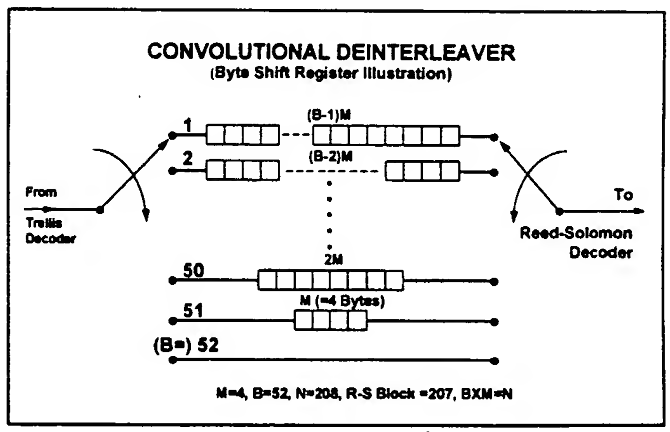

Reed-Solomon Decoder

The trellis-decoded byte data is sent to a (207,187) t=10 R-S decoder, which uses the 20 parity bytes to perform

byte-error correction on a segment-by-segment basis. Up to 10 byte errors per data segment are corrected by the R-S decoder. Any burst errors created by impulse noise, NTSC cochannel interference, or trellis-decoding errors, are greatly reduced by the combination of the interleaving and R-S error correction.

Data De-Randomizer

The de-randomizer accepts the error-corrected data bytes from the R-S decoder, and applies the same pseudo-random sequence (PRS) randomizing code to the data as was applied at the transmitter. The PRS code is generated using the same circuit (a shift register with the same feedback and output taps) as in the transmitter (Figure 18). Since the PRS is locked to the stable Data Frame Sync (and not some code word embedded within the potentially noisy data), it is exactly synchronized with the data, and performs reliably under all conditions.

Figure 18

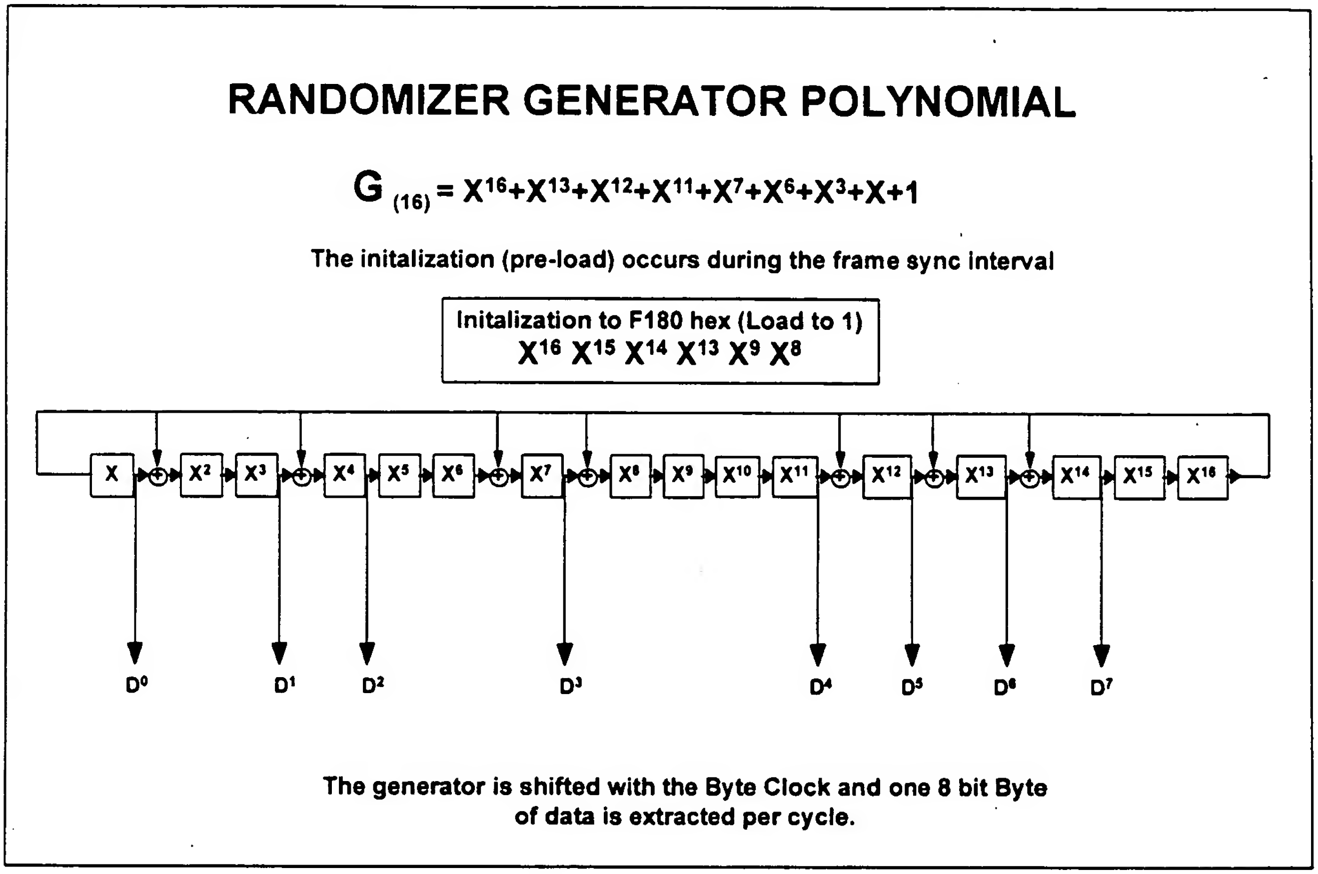

Sequence of Acquisition

The receiver takes advantage of the VSB signal features to first obtain robust lock to carrier, data symbols, and data framing information, all of which can be maintained through fades which destroy the data itself. Thus, there is no need to re-acquire the signal in a case where fading has momentarily caused loss of data.

The receiver incorporates a "universal reset" which initiates a number of "confidence counters" and "confidence flags" involved in the lock-up process. A universal reset occurs, for example, when tuning to a new station or turning on the receiver.

The various loops within the VSB receiver acquire and lock-up sequentially, with "earlier" loops being independent from "later" loops. The order of loop acquisition is as follows:

* Tuner 1st LO synthesizer acquisition
* Non-coherent AGC; Carrier (FPLL) acquisition (independent of each other)
* Data segment sync and clock acquisition

* Coherent AGC of signal
* Data frame sync acquisition
* - NTSC rejection filter insertion decision
* - Equalizer tap adjustments start
* - Data De-interleaver synchronization
* - Data de-randomizer synchronization
* Equalizer tap adjustments complete
* Phase Tracker acquisition
* Trellis and R-S data decoding start

Most of the loops mentioned above have confidence counters associated with them to insure proper operation. However, the build-up and let-down of confidence are not designed to be equal. The confidence counters build up confidence quickly for quick acquisition times, but lose confidence slowly to maintain operation in noisy environments. The VSB receiver carrier, sync and clock recovery circuits will work in S/N conditions of 0 dB or less as well as in severe interference situations.

High Speed Cable Mode Description

The high data rate cable mode trades off transmission robustness (28.3 dB signal-to-noise threshold) for system data rate (43 Mbit/sec). Most parts of the cable mode VSB system are identical or similar to the terrestrial system. As for the 8-VSB terrestrial signal, a pilot, data segment sync, and data frame sync are used to provide robust operation. The pilot in the cable mode adds 0.3 dB to the data power, as in the terrestrial mode. The symbol, segment, and frame signal rates are all the same in the two modes, allowing one type of receiver to lock up on either transmitted signal. Also, the data frame definitions are identical. The primary difference is the number of received levels (16 versus 8) and the omission of the trellis decoding and NTSC interference rejection filtering when receiving in cable mode.

The RF spectrum of the cable modem transmitter is identical to that of the terrestrial system, as illustrated in Figure 19. The error probability of 16-VSB, shown in Figure 20, is 3×10^{-6} BER at about 28.3 dB S/N, with forward error correction provided by Reed-Solomon coding. The error probability for terrestrial 8-VSB with its additional trellis coding also is shown in Figure 20 for comparison.

Figure 19

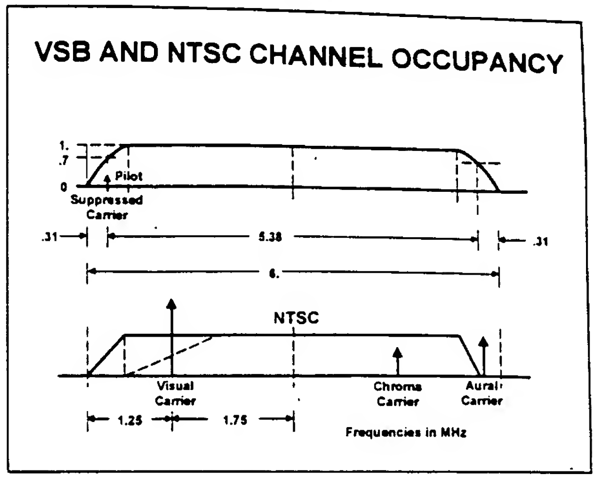

Figure 20

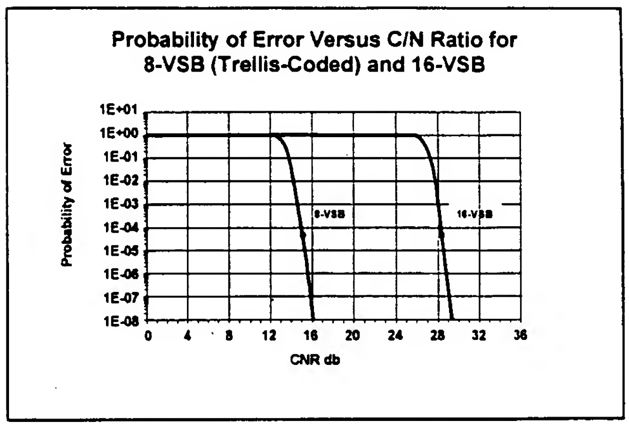

The cable-mode receiver, shown in Figure 21, is identical to the VSB terrestrial receiver, except that the trellis decoder is replaced by a slicer, which translates the multi-level symbols into data. Instead of an 8-level slicer, a 16-level slicer is used. Also note that no NTSC interference rejection filter is required in the cable mode since strong cochannel signals are not present on cable. The interleaver employed in the 16-VSB Cable mode is a 26 data segment, intersegment convolutional byte interleaver. Interleaving is provided to a depth of about 1/12 of a data frame (2 ms deep). The system will tolerate a burst error of 96.6 μsec.

Figure 21

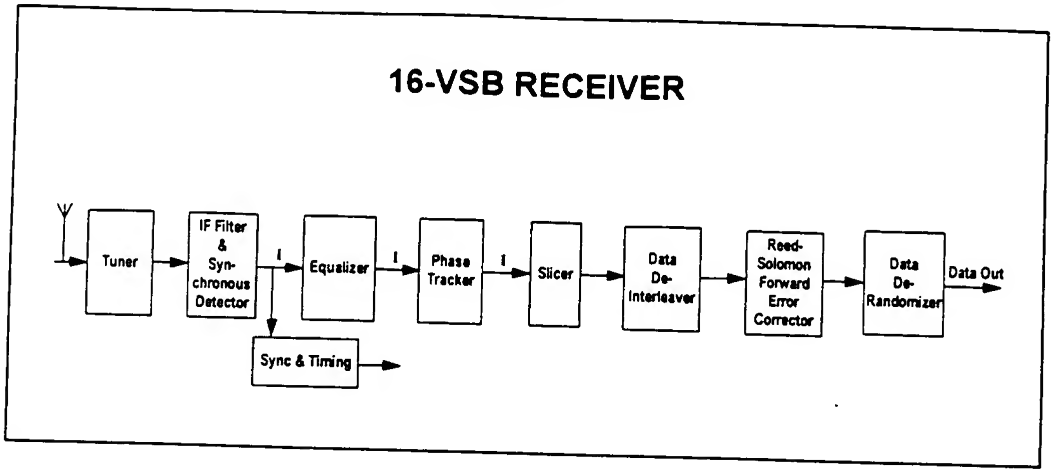

GENERAL REFERENCES

"Field Test Results of the Grand Alliance HDTV Transmission Subsystem", submitted to SS/WP2 Field Testing Task Force of the Advisory Committee on Advance Television Service of the Federal Communication Commission by Association for Maximum Service Television, Inc., Cable Television Laboratories, Inc., and Public Broadcasting Service, September 16, 1994.

"VSB Transmission System Grand Alliance: Technical Details", Monograph published by the Grand Alliance, December 7, 1994.

"Frequency and Phase-Locked Loop" by R. Citta, IEEE Transactions on Consumer Electronics, Volume CE 23, No. 3, August 1977, pp. 358-365.

ACKNOWLEDGMENTS

Thanks to the following members of the Electronic System Research & Development department at Zenith for their assistance in preparing this material: Carl Eilers, Gopalan Krishnamurthy, Brian Mehaffey, Dennis Mutzabaugh, Larry Nielsen, and Rudy Turner.

Wayne Bretl is a Staff Consulting Engineer in the Electronic Systems Systems Research and Development Department of Zenith Electronics. He received his BSEE from Illinois Institute of Technology in 1966. He holds 11 patents in television and related areas. He is a member of IEEE, SMPTE, and AES, and represents Zenith in several ATSC Technical Subgroups, other HDTV-related committees, and HDTV Grand Alliance Specialists Groups.

Gary Sgrignoli received his bachelor and master of science degrees from the University of Illinois, Champaign in 1975 and 1977, respectively. He joined Zenith Electronics Corporation in January 1977, and is currently a Senior Project Engineer in the Electronics Systems Research and Development Department. Mr. Sgrignoli has worked on television ghost canceling, cable TV scrambling, cable TV two-way data systems, and digital high definition transmission systems. He has published several technical papers, and holds 17 patents with 11 pending.

Paul Snopko is Manager of an Electronic Systems R&D group at Zenith Electronics Corporation, Glenview, Illinois. He received his B.S. and M.S.E.E. degrees from the University of Illinois, Urbana, in 1968 and 1969 respectively. Since joining Zenith full time in 1969 he has worked on a broad range of consumer electronics and related technologies including laser video disks, laser printing, thin film transistor LCD drive circuitry, manufacturing automation, telephone circuitry, power systems, display monitors, satellite TV, CATV, scan conversion and display systems for HDTV and compressed video and HDTV bitstream recording interfaces. His group is currently involved in VSB tuners and ICs and HDTV receiver development. Since 1969 he has been a member of the IEEE and since 1986 has served on the technical program committee of the ICCE which he chaired in 1992. He has been a member of the ATSC and Consumer Electronics Bus Committee of the EIA. He has been awarded 12 patents and has previously published in the Broadcast Technology, Consumer Electronics and Electron Devices Transactions of the IEEE and in the Journal of the SMPTE.